# LQCD-ext II
# CD1 Review

# Overview


## Paul Mackenzie
Fermilab
mackenzie@fnal.gov

For the USQCD Collaboration
http://www.usqcd.org


LQCD-ext II CD1
Germantown
February 25, 2014

# Synopsis

- The mission of the project is to acquire and operate dedicated hardware for the study of quantum chromodynamics (QCD).
- Hardware is located at BNL, FNAL and JLab.
- 2010-2014 LQCD Projects
  - LQCD-ext, (2010-2014) total budget $18.15 M for hardware and operations.
  - Sister project LQCD-ARRA, 2009-12, total budget $4.96 M.
  - Total 2010-2014 LQCD budget: $23.1M.
- Planned 2015-2019 LQCD-ext II budget: $14-18 M.

In light of the reduced budget envisioned for LQCD-ext II, we are making plans on the optimal way to slow down some of our deliverables without compromising our most critical goals.

Most of the talks in this review will be dedicated to project management and hardware strategy; this talk will discuss the scientific context in which this hardware will be used.

# In the past decade,

## USQCD has played a critical role in the HEP and NP experimental programs.

- In particle physics, lattice QCD supplied key hadronic matrix elements that made the Tevatron and flavor factory searches for physics beyond the standard model in heavy quark decays a success.
- In cold nuclear physics, lattice QCD supplied the exploratory calculations of the hadronic resonance and exotics spectra that guided planning for the JLab 12 GeV upgrade.
- In heavy-ion physics, lattice QCD supplied the chiral transition temperature that guided the interpretation of heavy-ion experiments like RHIC.

# In the coming decade,

## Even larger challenges are before us.

- HEP plans a new round of intensity frontier experiments, almost all of which require new kinds of lattice calculations;

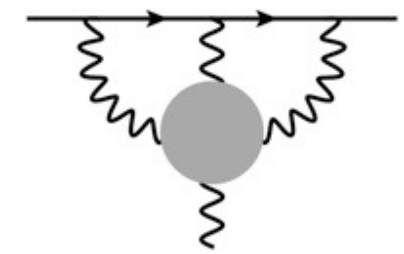

- The g-2 experiment needs lattice calculations of light by light scattering or it cannot succeed. Calculations this difficult have never yet been done.

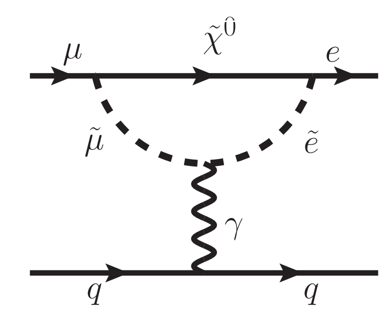


- The Mu2e experiment and proton decay experiments need nucleon matrix element calculations to interpret any beyond-the-standard-model signal it sees in terms of underlying physics models.

- Neutrino experiments like LBNE (and Nova, Minerva, ...) will need lattice calculations of nucleon matrix elements to sort out their murky mixture of nucleon and nuclear uncertainties.

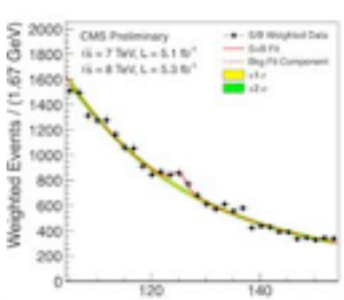


- High precision measurements of Higgs branching fractions will require high-precision lattice determinations of $m_b$ and $\alpha_s$ to eliminate parametric uncertainties in the search for beyond-the-standard-model effects.

- Nuclear physics (for example at FRIB) will require calculations of the light nuclei and their interactions at the physical pion mass to accomplish the systematic refinement of nuclear forces that underpin the interactions and structure of nuclei and the production of energy in stars.

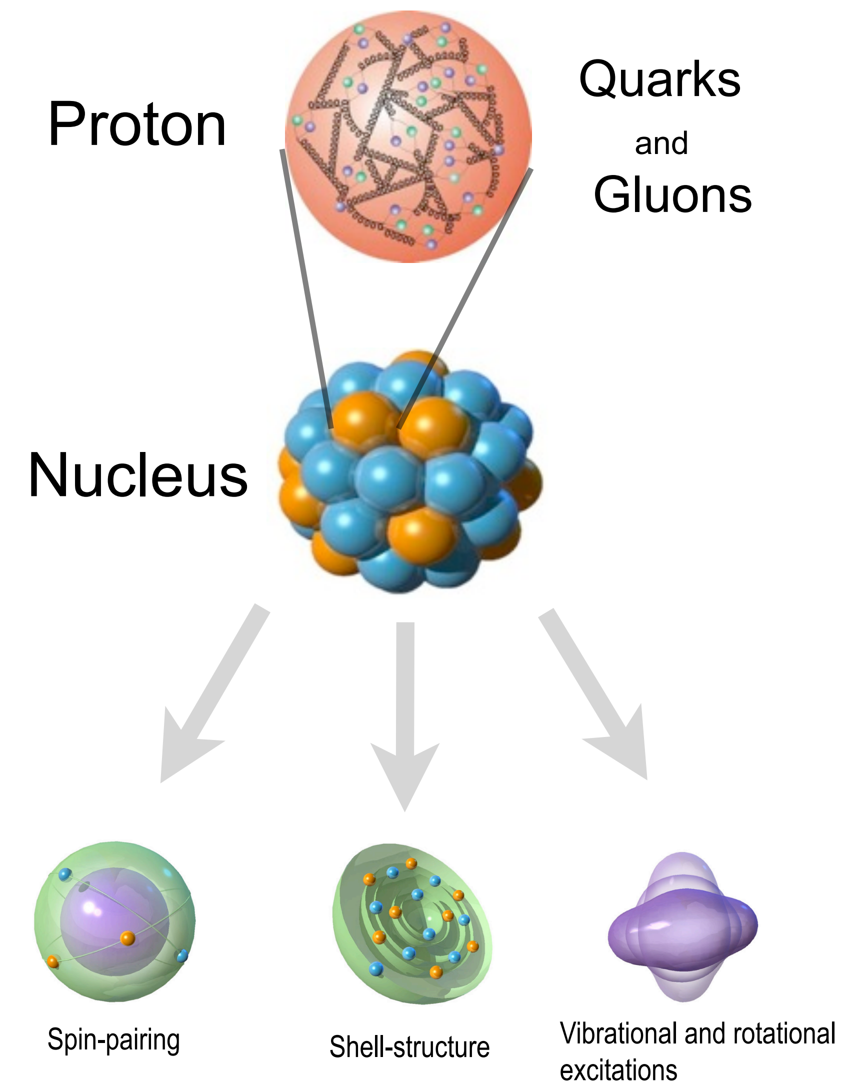


- The JLab 12-GeV upgrade will require solid determinations of excited hadronic masses and widths to compare with its experimental determinations.

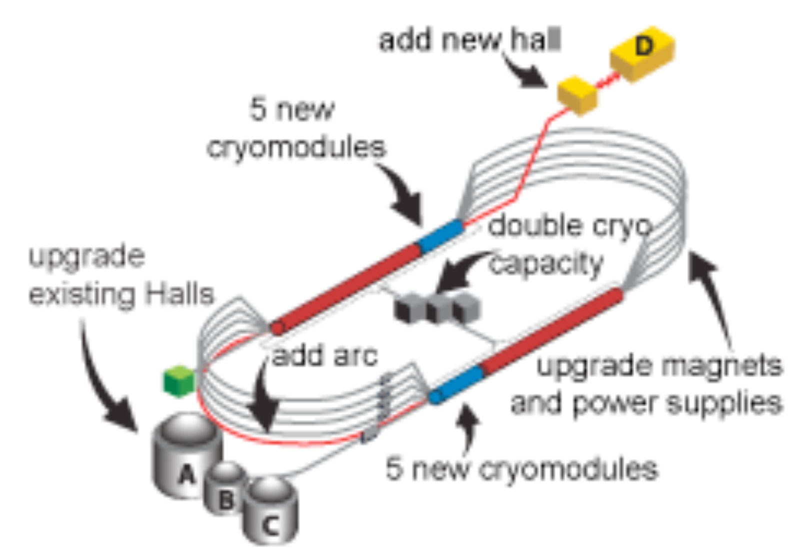

- Upcoming heavy ion experiments like the RHIC beam energy scan will need lattice QCD input on higher order cumulants of conserved charge fluctuations in order to be able to provide evidence for a possible critical point at non-zero baryon chemical potential.

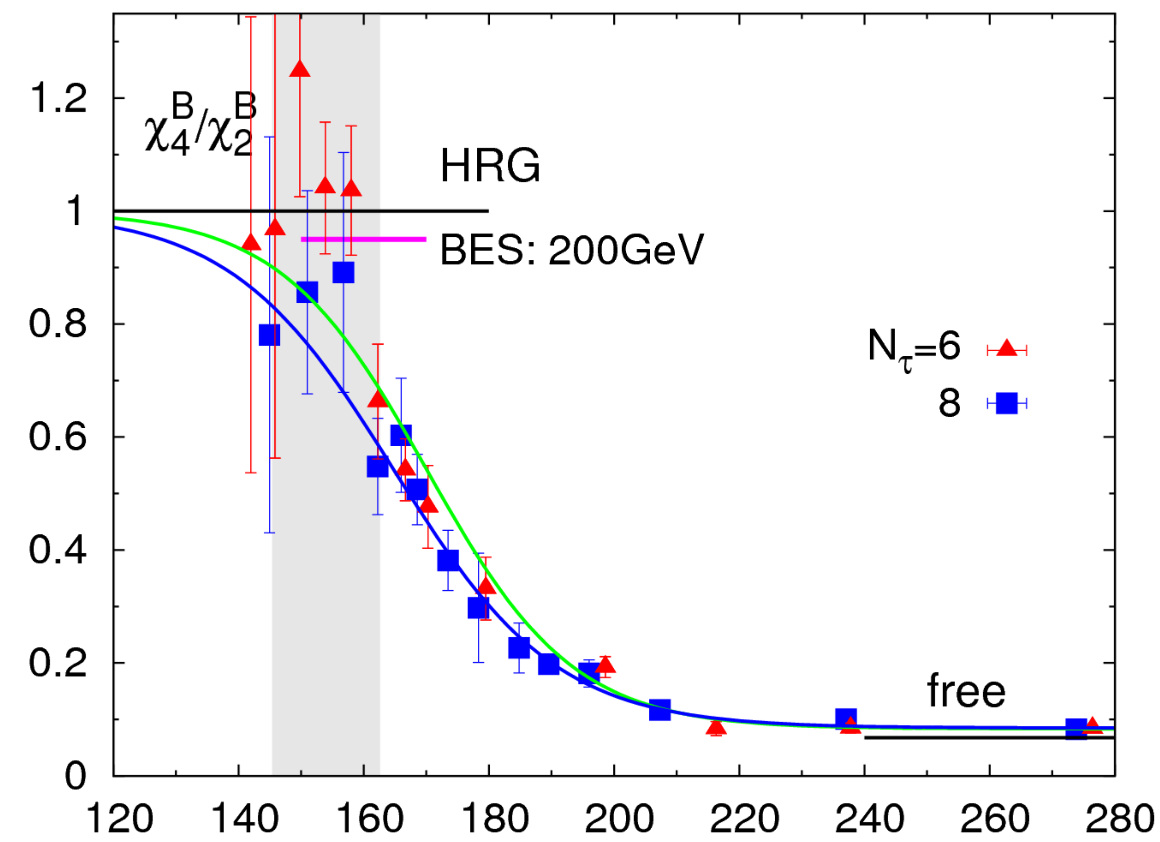

- In the coming five years, lattice gauge theory calculations have critical roles to play throughout the HEP and NP programs.
- As new experiments and facilities come online in HEP and NP, we can expect to find needs for lattice QCD in virtually all of them.
- The field of lattice QCD is not an undertaking which will go out of existence after accomplishing its goals.
  - It is a method of theoretical physics, like perturbation theory in QCD or QED, which will be needed for the foreseeable future.
- The LQCD Project is not a project like the Tevatron or the 12 GeV Upgrade, which will finish its mission and cease operations.
  - When the current round of calculations that it will enable are completed, new calculational needs will have arisen in their place.
  - It is currently funded in the US at a level of under 0.5% of the combined HEP and NP budgets, a critical level of support.

# The USQCD Collaboration

- Organizes computing hardware and software infrastructure for lattice gauge theory in the US,
  - including the LQCD Project and several other efforts.
- Represents almost all of the lattice gauge theorists in the US; ~ 163 people.
  - ~ 100 participating in physics proposals this year.
- Physics calculations are done by smaller component collaborations within USQCD:
  - Fermilab, HotQCD, HPQCD, HadSpec, LHPC, LSD, MILC, NPLQCD, RBC, ...

# USQCD main activities

- The LQCD Project
  - The subject of this review;  for highest capacity computing, requiring more floating point operations than are available at the Leadership Class Centers.
- Leadership Computing Centers
  - For highest-capability computing, that can’t be done on smaller computers.
  - DoE INCITE centers have been a mainstay; NSF Blue Waters beginning to play a role.
- DoE SciDAC grants
  - For software to run on all our computing resources.
  - Currently, ~$1M/year each from HP and HEP.

# Executive Committee

Paul Mackenzie (Fermilalb, chair), Rich Brower (Boston U.), Norman Christ (Columbia), Frithjof Karsch (BNL), Julius Kuti (UCSD), John Negele (MIT), David Richards (JLab), Martin Savage (U. Washington), Bob Sugar (UCSB)

- Provides overall leadership for the collaboration and point of contact for the DoE.
- Writes the proposals for hardware and software and chooses the members of the other committees.
- Rotates new members at ~ one/year.
  - Close to full rotation over ~ 10 years is planned. About half has rotated already.
  - We plan to rotate in a way that preserves rough balance between physics interests, HEP and NP, collaborations, etc.

# Scientific Program Committee

Robert Edwards (chair, JLab), Will Detmold (MIT), Anna Hasenfratz (Colorado), Taku Izubuchi (BNL), Peter Petretzky (BNL), Doug Toussaint (Arizona), Ruth Van de Water (Fermilab)

Each year, the many smaller physics collaborations within USQCD submit proposals to the Scientific Program Committee for allocations of time on USQCD's LQCD Project and Incite resources.

The SPC creates a program to accomplish the goals set forth in the USQCD Collaboration's proposals.

- It may also advise us on needed evolution of the goals.

The Executive Committee seeks the advice of the SPC on physics priorities when writing new proposals for DoE computing resources.

Chair rotates every two years. Members rotate every four years, at a rate of about two/year.

~30 people have served so far as members of either the EC or the SPC (almost 1/3 of the active members of the Collaboration).

# Goals: proposals and white papers

- The physics goals of USQCD are set out in our proposals and white papers by the Executive Committee in consultation with the SPC.
  - Continually evolving, in consultation with the SPC.
  - Discussed by the Collaboration at All Hands meetings.
- The LQCD-ext II proposal contains the most recent statement of our view of our most important goals and opportunities, and our view of our highest impact results. The 2013 white papers contain a more lengthy discussion. (See http://www.usqcd.org/collaboration.html .)
  - Outlined exciting physics program assuming approximate continuation of current support. ($23 M over five years from HEP and NP.)
- In setting and updating our goals we have always relied on informal input from numerous experimenters and phenomenologists.

# Science Advisory Board

- This year, we have formalized this process by naming a Science Advisory Board.
  - Brendan Casey (Fermilab, g-2), Marina Artuso (Syracuse, LHC-b), Jesse Thaler (MIT), David Kaplan (U. Washington), Curtis Meyer (Carnegie Mellon, GlueX), Nu Xu (LBL, Star), Volker Koch (LBL).
  - Among the most useful advisors on white papers and proposals.
- At the beginning of each year's allocation process, they will be asked to
  - Comment and suggest revisions of our general goals. (They have just done this for the first time.)
  - Read and comment on the year's physics proposals and allocations. (They will be asked to do this for the first time when this year's proposals are due on March 14.)
  - Invited to participate in meetings of SPC, in All Hands Meeting.
    - We are exploring how closely our advisors would like to be involved with the allocations process.

# Allocations

The jpsi core-hour is USQCD's standard allocation unit.
A jpsi core ~ 1.2 GF.
1 teraflop-year ~ 6.5 M jpsi ch.

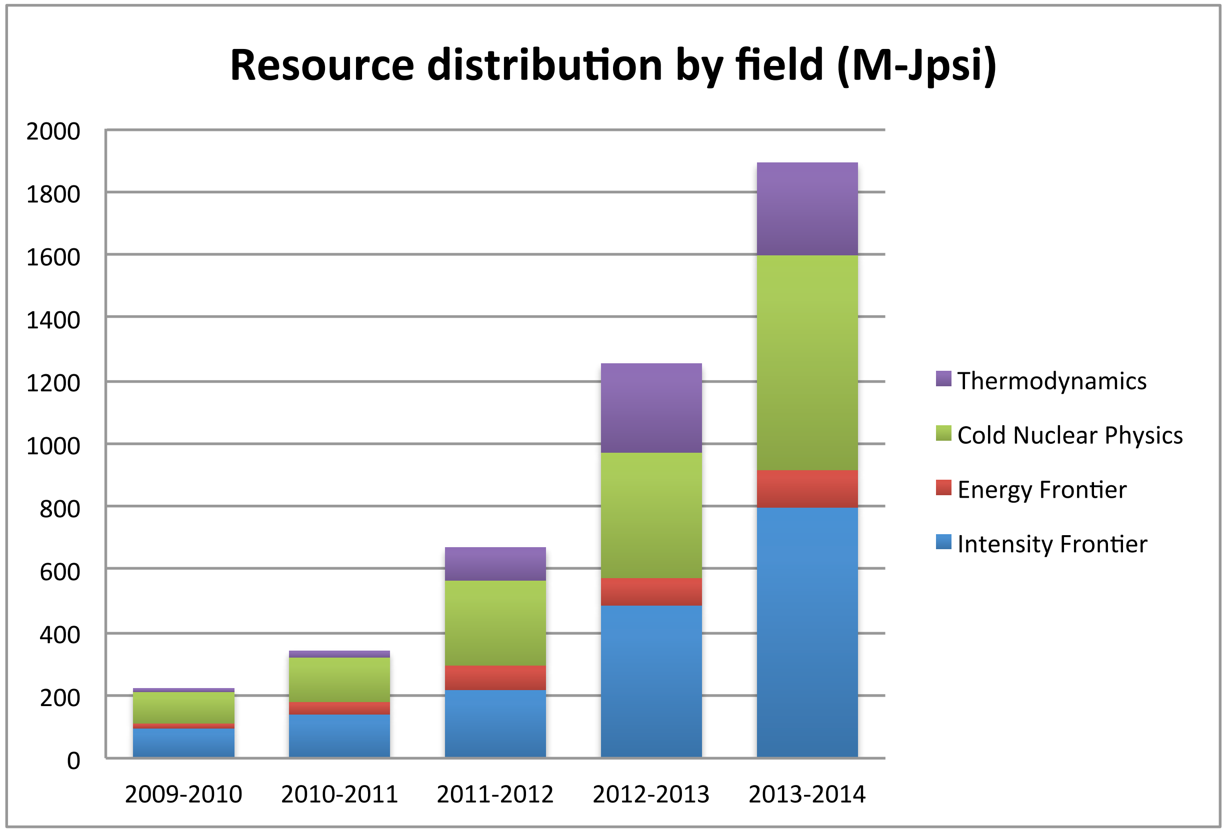


Beyond the standard model and QCD thermodynamics fractions have been rising.

GPUs 20% of Project $, but 50% of Project cycles; but: less general-purpose cycles:

- many GPUs are not error correcting, many single GPU jobs.

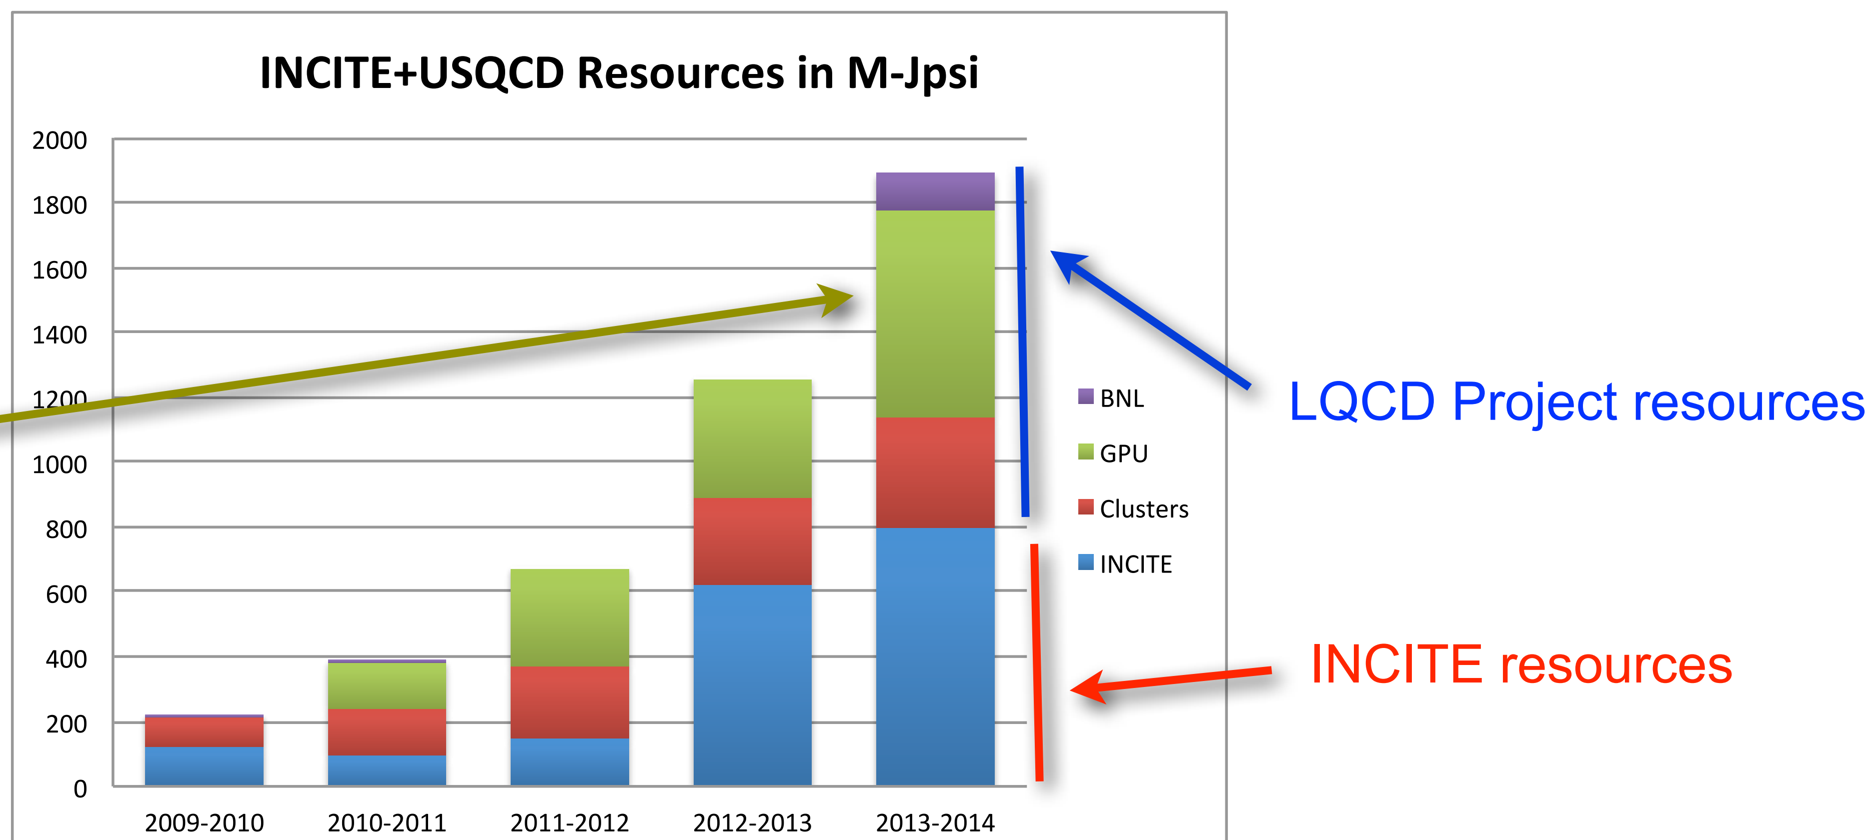

# Allocations

## Cluster allocations in 2013 by project.

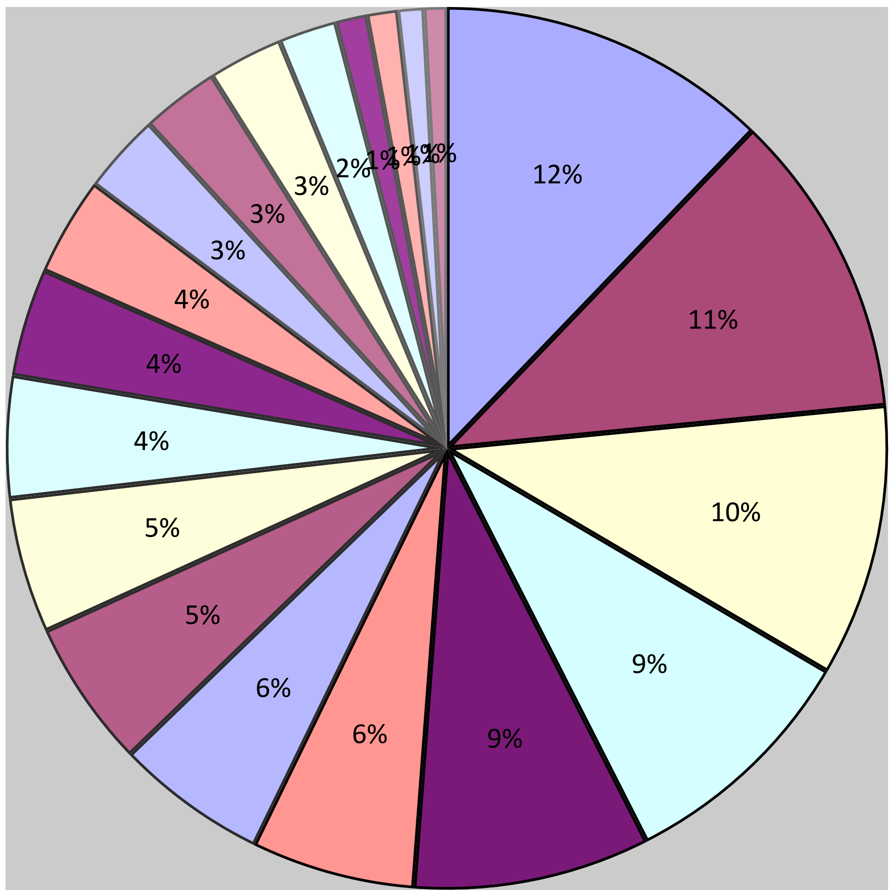


Projects are judged by:
- relevance to the central goals of USQCD;
- size and competence of project team;
- validity and efficiency of methods proposed.

Less high priority projects are typically not zeroed out, but are give less resources. About half of these allocations went to the five highest priority projects.

(Different from experimental programs, where experiments must be voted either up or down.)

The twenty smaller projects take up about half of our LQCD resources and, in the aggregate, are just as important to us as the biggest projects. Among other things, they contain the innovative and speculative calculations that are critical to the long-term future of our field.

# Anatomy of a typical lattice calculation

## Two main components:

100s of GB
file sizes

Generate *O*(1,000) gauge configurations on a leadership facility or supercomputer center. Hundreds of millions of core-hours.

Transfer to labs for analysis on clusters. Larger CPU requirements.

Gauge configuration generation:
a single highly optimized program,
very long single tasks,
"moderate" I/O and data storage.
Needs capability computing.

Hadron analysis.
Large, heterogeneous analysis code base,
10,000s of small, highly parallel tasks,
heavy I/O and data storage.
Needs capacity computing.

Two comparably sized jobs with quite different hardware requirements.

# Capacity and capability computing

- Leadership class computing is essential for generating large ensembles of gauge configurations. This computing cannot be done any other way.
- We have an even greater need for flops analyzing these configurations.
  - Can often be done very efficiently (cheaply in $/flop) in parallel on much smaller systems.

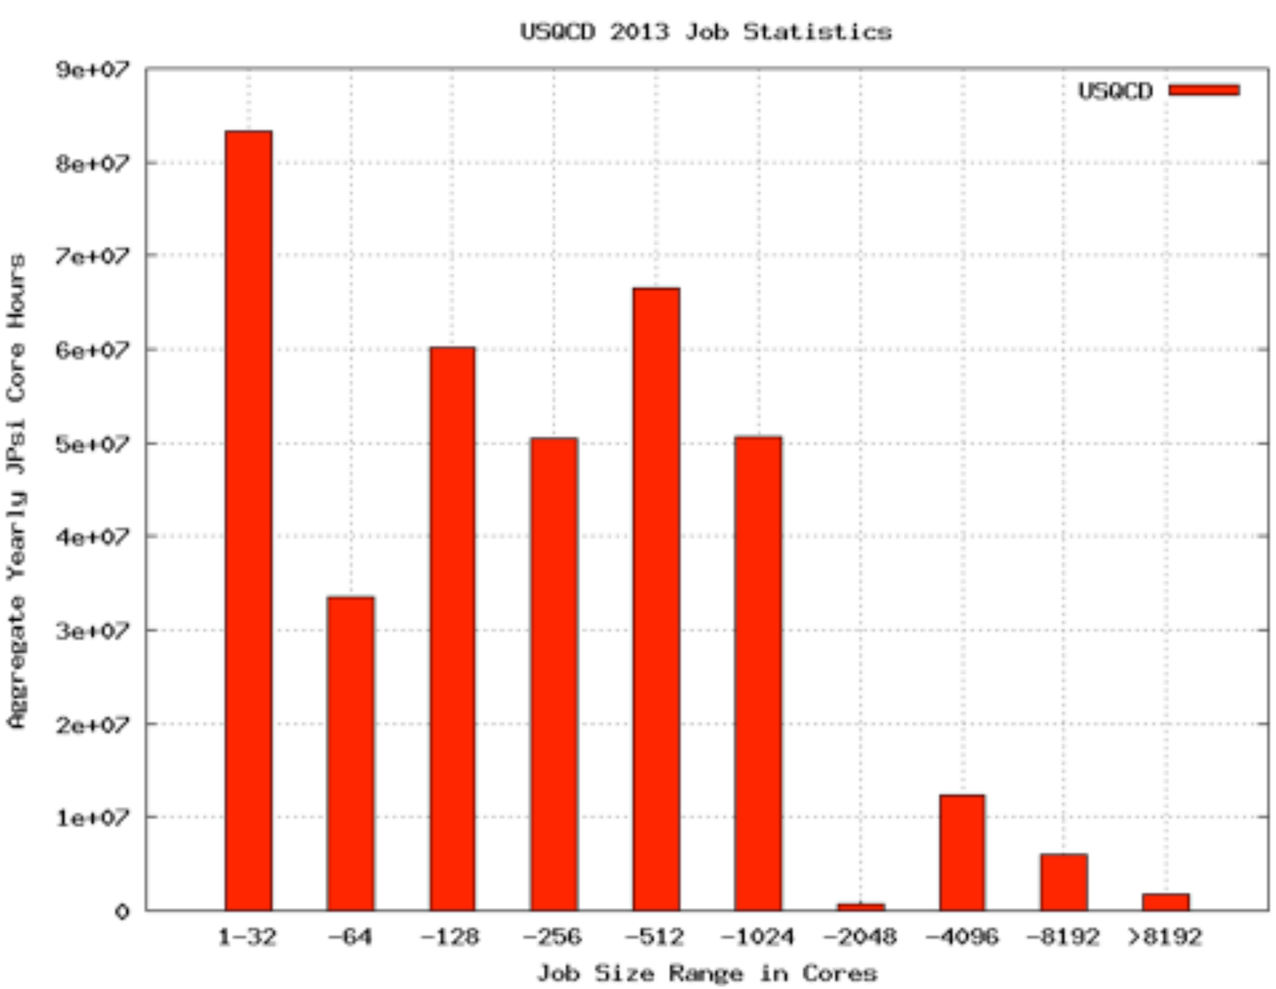


Job size distribution on USQCD 2013 conventional clusters.

We have an approximately flat distribution of job-size needs from one-node jobs to hundred thousand node jobs on a log scale in job-size.

The clusters of the LQCD Project supply the low end of these needs; the Leadership-Class Centers supply the high end.

# Roles for various kinds of hardware

- USQCD's conventional clusters and the Leadership Class facilities are the mainstays of USQCD's hardware.
- Some of our job mix, ~10-15%, can be run on single nodes. For these jobs, GPUs can offer far more flops/$ than conventional nodes.
    - JLab's 9g cluster has delivered this to great effect for one NP project.
- For projects that have overcome the programming challenge, larger GPU clusters at Fermilab and JLab supply more flops/$ than conventional hardware.
- For jobs in the 2K-8K range, USQCD operates a half-rack Blue Gene Q at Brookhaven.
    - The LCFs supply flops for jobs requiring>=16K cores well. GPU and conventional clusters supply flops for sizes <=1K cores well.

# Current resources

| | USQCD 2013 lattice QCD resources | | | | | |
|---|---|---|---|---|---|---|
| | Red: DoE resources | Subtotal | Core hours (M jpsi) allocated | Core hours (M jpsi) zero-priority | Comment | Policies |
| | | | | | | |
| **Lattice QCD** | **LQCD-ext, conventional.** | | 411 | | | |
| | **LQCD-ext, GPUs** | | 646 | | GPU-based* | |
| | | 1057 | | | | |
| | | | | | | |
| **Leadership-class** | **OLCF (DoE), Titan** | | 182 | | GPU-based | Jobs using < 20% of the machine ( nodes) don’t make it through the queue. |
| | **ALCF (DoE), Mira** | | 325 | 413 | | Proposed calculations using < 4-8 racks (64-128 K cores) are discouraged. 16 K core-jobs are allowed in zero-priority time. |
| | **ALCF Intrepid** | | 20 | 52 | (Being retired) | |
| | | 992 | | | | Total will fall in 2014 and (presumably) beyond. |
| | **NCSA(NSF), Blue Waters** | | 656 | | Partly GPU-based | No firm allocation policies announced yet. USQCD PRAC grant is entirely leadership-class, runs through April, 2015. Unknown whether part of a long-term program, and so is not part of USQCD long-term planning. |
| | | 656 | | | | |
| | | | | | | |
| **General purpose** | **NERSC (DoE)** | | 55 | | | |
| | **XSEDE (NSF)** | | 109 | | | |
| | | 164 | | | | |

*** The power of GPU-based systems is highly dependent on the calculations being run.**
**The rating of the LQCD GPUs was based on the particular mix of projects being run when the rating was done.**

# Current resources

| | Red: DoE resources | Subtotal | Core hours (M jpsi) allocated | Core hours (M jpsi) zero-priority | Comment | Policies |
|---|---|---|---|---|---|---|
| | USQCD 2013 lattice QCD resources | | | | | |
| | | | | | | |
| Lattice QCD | LQCD-ext, conventional. | | 411 | | | |
| | LQCD-ext, GPUs | | 646 | | GPU-based* | |
| | | 1057 | | | | |
| | | | | | | |
| Leadership-class | OLCF (DoE), Titan | | 182 | | GPU-based | Jobs using < 20% of the machine ( nodes) don't make it through the queue. OLCF |
| | ALCF (DoE), Mira | | 325 | 413 | | Proposed calculations using < 4-8 racks (64-128 K cores) are discouraged. 16 K core-jobs are allowed in zero-priority time. |
| | ALCF Intrepid | | 20 | 52 | (Being retired) | |
| | | 992 | | | | Total will fall in 2014 and (presumably) beyond. |
| | NCSA(NSF), Blue Waters | | 656 | | Partly GPU-based | No firm allocation policies announced yet. USQCD PRAC grant is entirely leadership-class, runs through April, 2015. Unknown whether part of a long-term program, and so is not part of USQCD long-term planning. |
| | | 656 | | | | |
| | | | | | | |
| General purpose | NERSC (DoE) | | 55 | | | |
| | XSEDE (NSF) | | 109 | | | |
| | | 164 | | | | |

*** The power of GPU-based systems is highly dependent on the calculations being run.**
**The rating of the LQCD GPUs was based on the particular mix of projects being run when the rating was done.**

# Current resources

| | USQCD 2013 lattice QCD resources | | | | | |
|---|---|---|---|---|---|---|
| | Red: DoE resources | Subtotal | Core hours (M jpsi) allocated | Core hours (M jpsi) zero-priority | Comment | Policies |
| | | | | | | |
| **Lattice QCD** | **LQCD-ext, conventional.** | | 411 | | | |
| | **LQCD-ext, GPUs** | | 646 | | GPU-based* | |
| | | 1057 | | | | |
| | | | | | | |
| **Leadership-class** | **OLCF (DoE), Titan** | | 182 | | GPU-based | Jobs using < 20% of the machine ( nodes) don't make it through the queue. OLCF |
| | **ALCF (DoE), Mira** | | 325 | 413 | | Proposed calculations using < 4-8 racks (64-128 K cores) are discouraged. 16 K core-jobs are allowed in zero-priority time. ALCF |
| | **ALCF Intrepid** | | 20 | 52 | (Being retired) | |
| | | 992 | | | | Total will fall in 2014 and (presumably) beyond. |
| | **NCSA(NSF), Blue Waters** | | 656 | | Partly GPU-based | No firm allocation policies announced yet. USQCD PRAC grant is entirely leadership-class, runs through April, 2015. Unknown whether part of a long-term program, and so is not part of USQCD long-term planning. |
| | | 656 | | | | |
| | | | | | | |
| **General purpose** | **NERSC (DoE)** | | 55 | | | |
| | **XSEDE (NSF)** | | 109 | | | |
| | | 164 | | | | |

*** The power of GPU-based systems is highly dependent on the calculations being run.**
**The rating of the LQCD GPUs was based on the particular mix of projects being run when the rating was done.**

# Current resources

| | | USQCD 2013 lattice QCD resources | | | | |
|---|---|---|---|---|---|---|
| | Red: DoE resources | Subtotal | Core hours (M jpsi) allocated | Core hours (M jpsi) zero-priority | Comment | Policies |
| | | | | | | |
| Lattice QCD | LQCD-ext, conventional. | | 411 | | | |
| | LQCD-ext, GPUs | | 646 | | GPU-based* | |
| | | 1057 | | | | |
| | | | | | | |
| Leadership-class | OLCF (DoE), Titan | | 182 | | GPU-based | Jobs using < 20% of the machine ( nodes) don't make it through the queue. OLCF |
| | ALCF (DoE), Mira | | 325 | 413 | | Proposed calculations using < 4-8 racks (64-128 K cores) are discouraged. 16 K core-jobs are allowed in zero-priority time. ALCF |
| | ALCF Intrepid | | 20 | 52 | (Being retired) | |
| | | 992 | | | | Total will fall in 2014 and (presumably) beyond. NCSA |
| | NCSA(NSF), Blue Waters | | 656 | | Partly GPU-based | No firm allocation policies announced yet. USQCD PRAC grant is entirely leadership-class, runs through April, 2015. Unknown whether part of a long-term program, and so is not part of USQCD long-term planning. |
| | | 656 | | | | |
| | | | | | | |
| General purpose | NERSC (DoE) | | 55 | | | |
| | XSEDE (NSF) | | 109 | | | |
| | | 164 | | | | |

* The power of GPU-based systems is highly dependent on the calculations being run.
The rating of the LQCD GPUs was based on the particular mix of projects being run when the rating was done.

# Policies of US supercomputing centers

- NERSC and XSEDE
  - General purpose computing. Member groups of USQCD apply for all the time they can get at these centers, but this time is dwarfed by Leadership-class resources and LQCD Project resources.
- Leadership class centers.
  - Aimed at computational tasks too large to be performed anywhere else.
  - OLCF ("Titan" GPU-equipped Cray): Aimed at tasks requiring >20% of the machine, smaller jobs don't make it through queue.
  - ALCF ('Mira" BG/Q): Aimed at jobs requiring 8 or more racks, smaller jobs receive deprecated allocations. (One-rack jobs may be run at zero-priority when allocation is finished.)
  - NCSA ("Blue Waters" GPU-equipped Cray) USQCD's allocation is purely leadership class. Allocation policies have not been announced.

# Could US program be run using Leadership-class centers only? No.

- A diverse program consisting of many projects ranging from small to very large is against the current policies of the LCFs.
  - USQCD program consists of ~30 projects varying widely in size.
- Even if the policies were changed, the LCFs are not set up with queue software or staff to monitor large numbers of small projects. (NERSC does that).
  - It would be unfeasible for USQCD to try to manage its entire program at an LCF using a single account.

- Even if an LCF and USQCD decided to run the entire USQCD program at the LCF, it would not be a cost-effective way for the DoE to supply our need for cycles.
    - Based on examination of LCF OMB300s and direct comparison of purchase price of leadership-class hardware. (One real world example carefully studied; >4x difference in price/performance.)

# Why are LQCD Project hardware resources so much more cost-effective than LCFs?

- The hardware is designed for the specific lattice calculations that are planned.
- Software and allocations are organized by the user community, not the responsibility of the Project.
- The USQCD user community contains many experts in high-performance computing:
  - Designers of the BG/Q,
  - Employees of the NVIDIA corporation,
  - Scientists who have run their own clusters.
- The Project expects users to be an aid to managing the system, not a burden needing support.

# Planned cost of LQCD-ext II

| LQCD-ext II Budget and Computing Capacity Scenarios | | | | | |
|---|---|---|---|---|---|
| **$18M Total Project Cost** | | | | | |
| | **Hdwr Budget** | **Ops Budget** | **Total Budget** | **New Deployments** | **Delivered** |
| **Fiscal Year** | **$M** | **$M** | **$M** | **TF** | **TF-years** |
| 2015 | 0.00 | 2.00 | 2.00 | 0 | 195 |
| 2016 | 1.61 | 2.39 | 4.00 | 107 | 185 |
| 2017 | 1.70 | 2.30 | 4.00 | 160 | 275 |
| 2018 | 1.84 | 2.16 | 4.00 | 244 | 475 |
| 2019 | 1.79 | 2.22 | 4.00 | 358 | 720 |
| Total | 6.94 | 11.06 | 18.00 | 869 | 1,850 |
| **$14M Total Project Cost** | | | | | |
| | **Hdwr Budget** | **Ops Budget** | **Total Budget** | **New Deployments** | **Delivered** |
| **Fiscal Year** | **$M** | **$M** | **$M** | **TF** | **TF-years** |
| 2015 | 0.00 | 2.00 | 2.00 | 0 | 195 |
| 2016 | 0.74 | 2.26 | 3.00 | 49 | 160 |
| 2017 | 0.99 | 2.01 | 3.00 | 93 | 190 |
| 2018 | 1.24 | 1.76 | 3.00 | 165 | 315 |
| 2019 | 1.33 | 1.67 | 3.00 | 250 | 480 |
| Total | 4.31 | 9.69 | 14.00 | 557 | 1,340 |
| **$23M Total Project Cost - Original (Proposal)** | | | | | |
| | **Hdwr Budget** | **Ops Budget** | **Total Budget** | **New Deployments** | **Delivered** |
| **Fiscal Year** | **$M** | **$M** | **$M** | **TF** | **TF-years** |
| 2015 | 2.63 | 1.85 | 4.48 | 165 | 270 |
| 2016 | 2.63 | 2.02 | 4.65 | 233 | 410 |
| 2017 | 2.63 | 2.07 | 4.70 | 330 | 650 |
| 2018 | 2.63 | 2.13 | 4.76 | 467 | 1,040 |
| 2019 | 2.63 | 2.18 | 4.81 | 660 | 1,560 |
| Total | 13.15 | 10.25 | 23.40 | 1,855 | 3,930 |

Current cost planning, high end.

Current cost planning, low end.

Planning in original proposal.
(More optimistic assumptions than used above.)

The $14M low end of the current planning guidance will allow us to operate the existing hardware, but it will very significantly curtail the anticipated growth in computing power.

- The significantly reduced funding levels and computing power now planned will significantly constrain our physics aspirations.
  - The collaboration and the Project team are looking at ways of reducing the planned budget for operations and at ways of adjusting the schedule for physics deliverables.
- The Executive Committee and the SPC are working to plan the optimal way to slow down some of our deliverables without compromising our most critical goals.
  - In discussion now, and will be refined in the coming year.
  - The final plans will be determined on a year-by-year basis by the SPC in consultation with the Executive Committee. The following are some of the issues we will be facing then and some of the possibilities we will have to consider.

## Core areas in the HEP lattice program:

- Previously, in the Tevatron/flavor-factory era, flavor physics and the parameters of the Standard Model were the core of the program.
- As the LHC has come on, small but increasing amounts of effort have gone into new strongly interacting gauge theories which may appear at the higher energy LHC
- The coming US experimental program will have new needs: g-2 first and foremost, plus Mu2e, neutrino nucleon matrix elements, proton decay matrix elements, etc.

Core areas in the HEP lattice program:

- Previously, in the Tevatron/flavor-factory era, flavor physics and the parameters of the Standard Model were the core of the program.
- As the LHC has come on, small but increasing amounts of effort have gone into new strongly interacting gauge theories which may appear at the higher energy LHC
- The coming US experimental program will have new needs: g-2 first and foremost, plus Mu2e, neutrino nucleon matrix elements, proton decay matrix elements, etc.

Regardless of funding, our plan has been to slow down the still important flavor calculations in favor of those needed by the future experimental program.

# Core areas in the HEP lattice program:

With the significantly reduced flops now envisioned, we will try to protect the calculations needed for experiments coming on line as much as possible (such as g-2 in 2016). We will probably have to delay other important calculations.

The flavor physics calculations still needing to be done will probably have to be further slowed down, delaying the discovery of possible BSM effects that could be found with the improved calculations.

Until and unless the LHC discovers new physics, investigation of possible new strongly coupled models for new physics will probably have to be slowed down. This is risky, because the 14-TeV LHC could discover new physics any time after it starts in 2015, leaving us less prepared at the outset.

Core areas in cold NP lattice program:

- Hadron resonance spectroscopy - the nature of gluon fields within hadrons, needed for GlueX experiment in 2015/16.

Critical to an experiment coming online - try to protect?

- Nucleon structure - the origin of spin, needed for 12-GeV upgrade: generalized parton distributions, fragmentation functions, form factors.

Critical to an experiment coming online - try to protect?

- Nuclear structure - origin of nucei from QCD - and the production of energy in stars.

A fundamental question in NP - the foundation of all of nuclear physics - but a longer-term project. May need to be slowed down, with prototype calculations continuing but definitive results pushed farther into the future.

Core area in heavy-ion collision lattice program:

- QCD thermodynamics calculations - matter at extreme temperatures and densities, needed to analyze RHIC physics, such as the upcoming RHIC Beam Energy Scan.

Cumulants to 4th order for fluctuation measurements in BES-II are of highest timeliness. >=6th order cumulants needed for better constraints on the critical point may be delayed.

Heavy quark spectral functions and diffusion constants are needed soon in the analysis of heavy quark spectroscopy, but precision calculations of light quark spectral functions needed for dilepton and photon rates, electric conductivity may be delayed

As the final scientific program is decided by the SPC and the Executive Committee each year in 2015-19, these are some of the issues and possibilities we will have to be considering.

# Summary

- The LQCD Project has made possible USQCD's large, diverse physics program, including results critical to the success of the HEP and NP experimental programs.
  - Provides efficient, medium-scale computing for the most important, flagship projects;
  - Supports broad, diverse program of smaller projects for innovation;
  - Makes possible serious physics projects by post-docs and young people.
- NP and HEP are planning an exciting future experimental program which will require new lattice calculations. The LQCD-ext II is poised to make these possible.
- A budget at the lower end of our current guidance is challenging our ability to deliver the calculations critical to the future programs of HEP and NP. We are working to find the best optimization possible.

# Backup slides

# SciDAC lattice QCD software R&D

The third critical component of our computational infrastructure.

Software Committee: Richard Brower (chair).

Regular Thursday phone conferences for people working on USQCD software.

USQCD has SciDAC-3 grants from HEP and NP for about $1 M each for creating lattice QCD software infrastructure: community libraries, community codes, optimization and porting to new architectures, implementation of up-to-the-minute algorithm advances...

- The QCD API and community libraries
  - Lower entrance barriers to lattice QCD.
  - Enable postdocs to run major projects without being part of major collaborations.
- Porting and optimizations for new platforms
  - Critical to efficient use of new hardware.

# World lattice computing resources

| Country | Sustained TF/sec | M jpsi ch/year |
| --- | --- | --- |
| **Germany** | 390 | 2,535 |
| **Japan** | 260 | 1,690 |
| **United Kingdom** | 260 | 1,690 |
| **United States** | | |
| **LQCD Project** | 163 | 1,057 |
| **DoE Leadership Class Centers** | 153 | 992 |
| **US Total** | 315 | 2,049 |
| | | |

Major resources in sustained teraflops/s estimated to be available for the study of lattice QCD in various countries in 2013.
1.0 TF-year = 6.5 M jpsi core-hours.

USQCD's computing resources put it in the big leagues, but they by no means dominate.

# Lattice QCD meets experiment meetings

To increase the interaction between lattice gauge theory and experiment and phenomenology, members of USQCD have organized a series of workshops with experimenters and phenomenologists.

- SLAC, Sept. 16, 2006, Standard Model physics.  With BaBar.
- Fermilab, December 10-11, 2007, “Lattice Meets Experiment” in flavor physics.
- Livermore, May 2-3, 2008, “Lattice Gauge Theory for LHC Physics”.
- JLab, Nov. 21-22, 2008, “Revealing the Structure of Hadrons”, Nuclear.
- BNL, June 8-9, 2009, “Critical Point and Onset of Deconfinement”, QCD thermodynamics.
- BU, Nov. 6-7, 2009, “Lattice Gauge Theory for LHC Physics”.  BSM.
- Fermilab, April 26-27, 2010, “Lattice Meets Experiment” in flavor physics.
- BU, 8-10 September 2010, “Sixth Workshop on QCD Numerical Analysis, Boston.
- JLab, Feb. 23-25, 2011, “Excited Hadronic States and the Deconfinement Transition”.
- BNL, Oct. 3-5, 2011, "Fluctuations, Correlations and RHIC low energy runs".
- Fermilab, Oct. 14-15, 2011, “Lattice Meets Experiment: Beyond the Standard Model”.
- Boulder, Oct 28, 2012, “Lattice Meets Experiment 2012: Beyond the Standard Model”.
- George Washington University, Aug. 21-23, 2012, “Extreme QCD”.

# DoE computing resources planned for lattice QCD, 2015-2019

| Fiscal Year | Dedicated Hardware (TF-Years) | Leadership Class Machines (TF-Years) |
|---|---|---|
| 2015 | 325 | 430 |
| 2016 | 520 | 680 |
| 2017 | 800 | 1,080 |
| 2018 | 1,275 | 1,715 |
| 2019 | 1,900 | 2,720 |
| Total | 4,820 | 6,625 |

LQCD-ext II Proposal

Currently, it appears that GPU resources will supply a larger fraction of our needs than was planned in the proposal, delivering a higher number of flops.
Both DoE LCFs are now expected to upgrade at the same time. The due date of the next generation of LCF machines (the “CORAL” consortium machines) is slipping, currently in 2017 and it appears that they will supply fewer cycles than we had hoped.
The aggregate of the expected DoE resources continues to be about as planned.

# 2013 USQCD all-hands meeting

- Took place April 19-20, 2013 at BNL.  ~80 members attended.  (http://www.usqcd.org/meetings/allHands2013/.)
- Reports from the Executive Committee, the LQCD-ext Project Manager, the SPC, and the hardware site managers.

- In each science domain, reports from
  - white paper authors,
  - representative physics projects,
  - members of the SPC on the relation between the allocated projects and the long-term goals.
- Round table discussions on
  - Coherence of the USQCD Software, Hardware and Physics Programs,
  - Evolution of USQCD Scientific Goals, Proposal Process, and Computational Priorities

# Membership survey and demographic information

- We are starting to collect membership and demographic information in a more organized way.
  - New membership list. Currently, ~ 163 members.
  - Demographic survey.

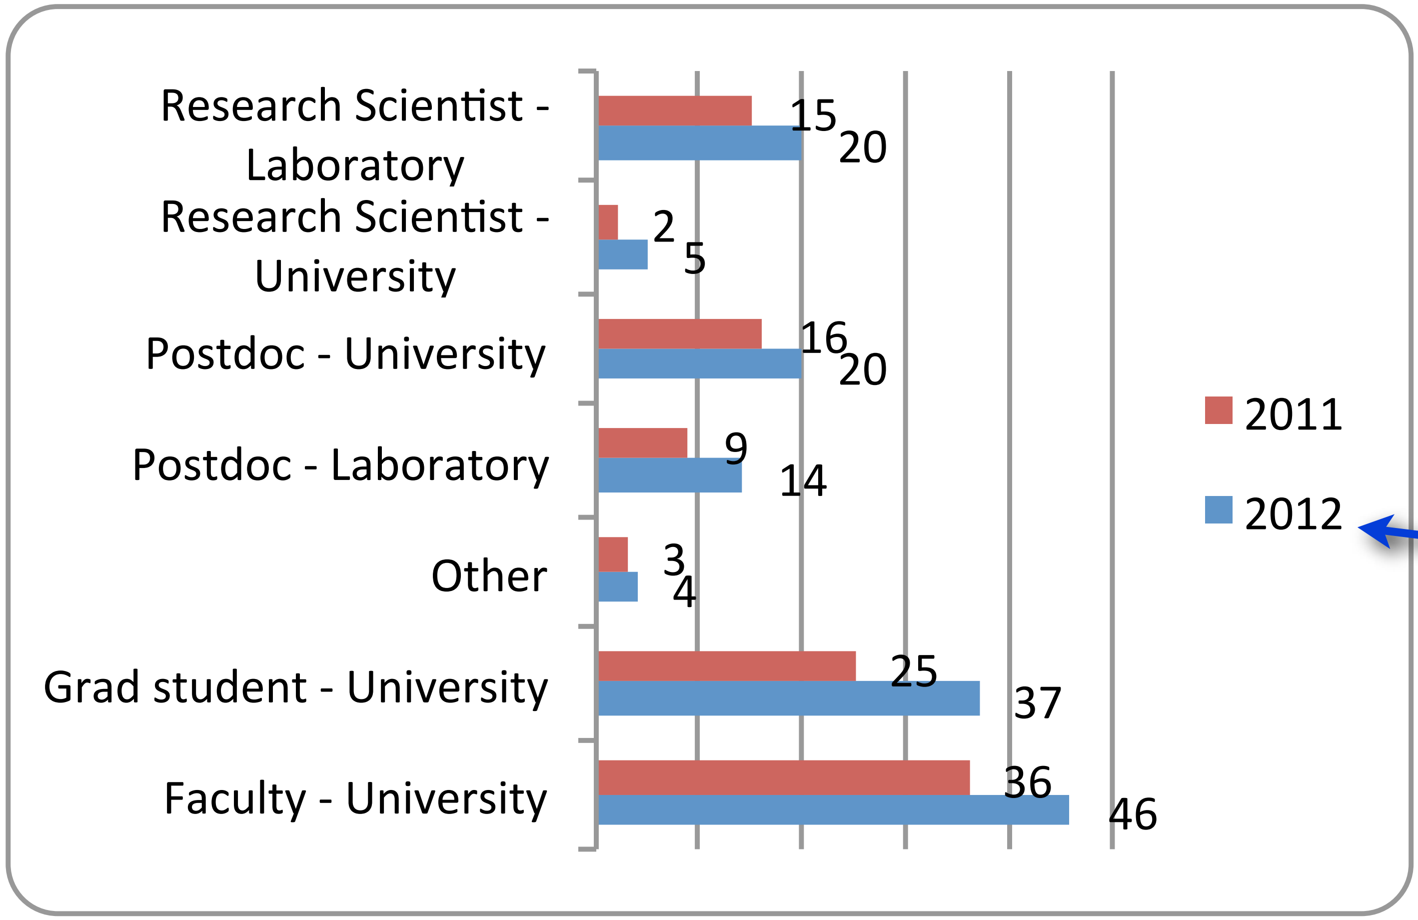


Better response to the survey in 2012, not a increase in the field.

Intend to make it annual.

We've grown from about 90 people in 2000 to about 163 today.

# International collaboration

- Lattice QCD is an international field with very strong programs in Germany, Italy, Japan and the United Kingdom, and elsewhere. Groups within USQCD have formed a number of international collaborations:
  - The USQCD effort using DWF quarks is an international effort between the United States based RBC, the Edinburgh, and Southampton members of the UKQCD Collaboration, and RIKEN.
  - The Fermilab Lattice, HPQCD and MILC Collaborations have worked together in various combinations to study heavy quark physics using improved staggered quarks. HPQCD includes physicists in both USQCD and UKQCD.
  - Members of the BNL Nuclear Physics lattice gauge theory group have a long term collaboration with physicists at the University of Bielefeld, Germany.
  - Members of USQCD working on the hadron spectrum using Clover quarks on anisotropic lattices have close ties with colleagues in Trinity College, Dublin, the Tata Institute, Mumbai, Cambridge U.
  - ...

# Junior staff job creation

- In the last few years, job creation has been good. About 25 permanent jobs in the US in the last 10 years; 30 or more would be a more comfortable number.
- We are working on strategies to improve it.
  - We have formed a speakers committee to find prominent speaking slots for talented young people.
    - Andreas Kronfeld (chair), David Richards, Peter Petreczky, Simon Catterall.
  - NP has had success creating university jobs with JLab bridge positions. We investigated the possibility for HEP.
    - DoE: DoE theory supports this idea in principle. They asked us to bring it to them again when the research budget stabilizes.
    - NSF: NSF avoids pushing universities in any particular direction.
    - Riken: the Japanese research organization Riken has agreed to create a joint job at the Riken BNL center joint with U of Colorado,
      - Went to Ethan Neil. BNL is invesitgating the possibility of more positions.

# Junior faculty and staff job creation

| | Year | Research institution, HEP | Research institution, NP | Computational scientist | Teaching college | Industry | Foreign |
|---|---|---|---|---|---|---|---|
| **Jack Laiho** | **2013** | Syracuse | | | | | |
| **Will Detmold** ** | **2013** | | MIT | | | | |
| **Ethan Neil** | **2013** | Colorado | | | | | |
| **Christopher Thomas** | **2013** | | | | | | Cambridge |
| **Ruth Van de Water** | **2012** | Fermilab | | | | | |
| **Elizabeth Freeland** | **2011** | | | | Benedictine U. | | |
| **Brian Tiburzi** | **2011** | | CUNY | | | | |
| **Andrei Alexandru** * | **2011** | | GWU | | | | |
| **Elvira Gamiz** | **2011** | | | | | | Granada |
| **Mike Clark** | **2011** | | | | | NVIDIA | |
| **Ron Babich** | **2011** | | | | | NVIDIA | |
| **Christopher Aubin** | **2010** | | | | Fordham | | |
| **Swagato Mukherjee** | **2010** | | BNL | | | | |
| **Changhoan Kim** | **2010** | | | | | IBM | |
| **Enno Scholz** | **2009** | | | | | | Regensburg |
| **Taku Izubuchi** | **2008** | BNL | | | | | |
| **James Osborn** | **2008** | | | Argonne | | | |
| **Chris Dawson** | **2007** | Virginia | | | | | |
| **Nilmani Mathur** | **2007** | | | | | | Tata Institute |
| **Joel Giedt** | **2007** | RPI | | | | | |
| **Matthew Wingate** | **2006** | | | | | | Cambridge |
| **Jozef Dudek** ** | **2006** | | Old Dominion | | | | |
| **Jimmy Juge** | **2006** | | | | U. of the Pacific | | |
| **Peter Petreczky** | **2006** | | BNL | | | | |
| **Balint Joo** | **2006** | | | JLab | | | |
| **Kieran Holland** | **2006** | | | | U. of the Pacific | | |
| **Kostas Orginos** ** | **2005** | | Wm & Mary | | | | |
| **George Fleming** | **2005** | | | Yale | | | |
| **Tom Blum** * | **2003** | Connecticut | | | | | |
| **Silas Beane** * | **2003** | | UNH | | | | |
| | | | | | | | |
| **Total** | | 7 | 8 | 3 | 4 | 3 | 6 |

* NSF Early Career Award
** DoE OJI/Early Career

# The LQCD Project is valuable to all its stakeholders.

- It’s good for the offices of HEP and NP because it is essential for meeting the science mission needs of the offices.
- It’s good for the leadership computing centers because it supplies the essential analysis cycles required to make the ensembles generated at the LCFs, valuable; capacity cycles that the LCFs are mandated *not* to supply.
- It’s good for the laboratories because it helps accomplish the lab experimental programs, and it provides design ideas and prototyping that are useful to other lab programs, e.g., prototyping GPUs and Intel PHI chips.

# Capacity and capability computing

- The LCFs are mandated to supply only capability computing needs, jobs that can't be done on any other machines.
    - >128,000 core jobs at ALCF, >3,600 GPU jobs at OLCF.
- Lattice QCD has an even greater need for capacity cycles that LCFs are mandated *not* to supply.
    - The LCFs are well aware that this need exists and are very happy that it is being met because it gives value to the large ensembles we generate there.
- This capacity need can be supplied much more efficiently on dedicated lab clusters than at multipurpose computing facilities.
    - Compare 200 M jpsi core-hours (140 M Oak Ridge core hours) for $7M at LCF vs >400 M jpsi core-hours for an average of $4.9 M/year on clusters at a HEP or NP lab - >3X more cost effective.

# HEP and NP labs are well suited to supply this need.

- LQCD clusters leverage lab capabilities.
  - Cluster expertise (reconstruction farms and real time triggers)
  - Storage (networks, file systems, data movement)
- LQCD hardware often provides design ideas and prototyping that is useful to other programs at labs and universities.
  - E.g., at Fermilab, for several years Ds-type machines (quad-socket Opterons) have been the standard used for Run 2, FermiGrid, CMS Tier 1;
  - other programs at the labs are now becoming very interested in GPUs and Intel MIC architecture.

# GPUs

- GPUs have supplied a significant fraction of our capacity computing needs in the last few years.
  - A disruptive technology: for the projects that can use them, they are very price performant in $, but require significant investments in software and physics brain power.
  - Price performance varies much more by project than is true for ordinary clusters.
    - Harder to define a standard candle for price/performance.
  - Knowledge of which projects can use them well is incomplete.
    - Harder for the Executive Committee to advise the project team about foreseen use.
  - The speedup we find has been enabled by the terrific work of our software committee.
  - ARRA funding at JLab was very valuable because it let USQCD prototype and then build GPU clusters at scale, accelerating software development and uptake.

# Organization, goals, allocations

- In 2003 when USQCD hardware funding began, Peter Rosen (head of HEP and NP) made it clear that DoE expected the hardware to be operated as a national facility.
    - Open to all in US to submit proposals.
    - Like Fermilab; not like BaBar.
    - Overall physics goals are set by USQCD in our white papers and proposals for hardware and software, but specific projects are developed by component collaborations like MILC, RBC, NPLQCD, HOTQCD, ..., and allocated by SPC. (Role of EC in this process is analogous to that of lab director.)
- We think this model has worked very well.
- A different model: USQCD could function as a physics collaboration like BaBar.
    - Not open to all in US; individuals would apply to join and specify their contribution in advance.

- Advantages of the facility model.
  - Young people can be PIs of their own physics programs as soon as they are able to formulate a project and a proposal that is convincing to the Scientific Program Committee.
    - They can be recognized for their own scientific programs much more easily than as part of a hundred-member collaboration.
    - The four people who got junior faculty jobs this year all served as PIs of their own proposals; two of them with no senior collaborators.
  - When groups disagree on methods, they can compete.
- Possible advantages of the collaboration model?
  - Could more straightforwardly enforce that the most important goals are implemented in the allocation process? We think that the current allocation process does this well.
- We think that the advantages outweigh the disadvantages of the facility model.

# US lattice gauge theory work flow

Zero-temperature QCD:

Currently three main streams of QCD gauge configurations are being generated by USQCD for different physics goals:

These high-value ensembles are data-rich resources that are shared among all of USQCD.

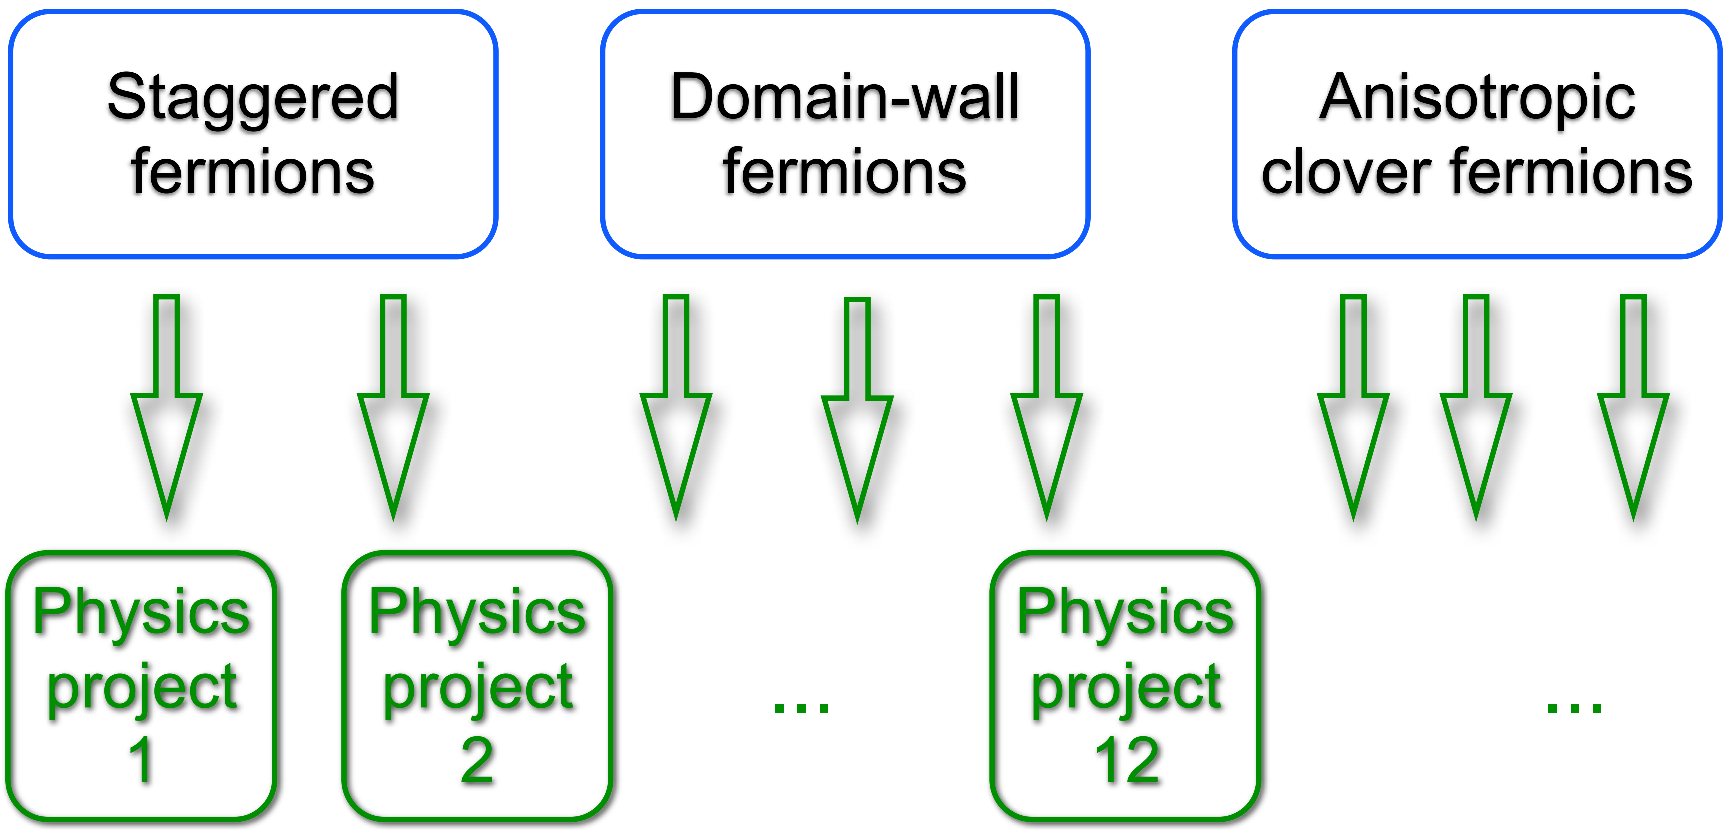


Shared among a couple of dozen groups, in both HEP and NP.

Physics projects are done on these configurations by smaller groups of 5-15 members within USQCD.

Around 90 of the 163 members of USQCD have submitted jobs to USQCD hardware.

QCD thermodynamics and BSM projects generate their own configurations tailored to specific goals.

# 2012 USQCD physics projects

Projects for configuration generation

Non-QCD Beyond-the-Standard-Model projects.

(Some QCD calculations are also looking for BSM effects in hadrons.)

Projects initiated by post-docs or students

**2012 Type A Proposals (more than 2,500,000 J/Psi core-hours)**

| PI | Project |
|---|---|
| Christopher Aubin | Hadronic contributions to the muon g–2 using Asqtad staggered fermions |
| Alexei Bazavov | Equation of State for 2+1 flavor QCD using Highly Improved Staggered Quarks |
| Thomas DeGrand | Running coupling and anomalous dimension of SU(3) gauge theory with adjoint fermions |
| Carleton DeTar | Quarkonium Physics in Full QCD |
| William Detmold | Lattice QCD Study of Multi-Baryon Systems |
| Heng-Tong Ding | Universal properties of the chiral phase transition in 2+1 flavor QCD using Highly Improved Staggered Quark action |
| Robert Edwards | Dynamical Anisotropic-Clover Lattice Production |
| George Fleming | Lattice Gauge Theory for Physics beyond the Standard Model on Leadership Class Machines |
| Anna Hasenfratz | Many-flavor gauge theories: Finite volume scaling at small masses |
| Taku Izubuchi | Applications of QCD+QED simulations (II): Isospin breaking in the hadron spectrum and Hadronic contributions to the muon anomalous magnetic moment |
| Julius Kuti | Lattice search for the BSM composite Higgs mechanism and the conformal window |
| Huey-Wen Lin | Probing TeV Physics through Neutron-Decay Matrix Elements |
| Keh-Fei Liu | Nucleon Form Factors and Hadron Spectroscopy |
| Paul Mackenzie | CKM Physics from B, D, and K Mesons |
| Robert Mawhinney | Pion and Kaon Physics from 2+1 avor DWF Lattices with m_pi = 140 MeV and V=(5.5 fm)^3 |
| Swagato Mukher | Continuum limit of higher-order charge-fluctuations at the physical point |
| John Negele | Precision Calculations to Extract Nucleon Ground State Structure in the Chiral Regime |
| Ethan Neil | Extended Study of Many-Fermion Gauge Theories for TeV Physics |
| Kostas Orginos | Isotropic Clover Fermions |
| David Richards | Excited Meson and Baryon States using Anisotropic Clover Lattices |
| Junko Shigemitsu | High-Precision Heavy-Quark Physics |
| Eigo Shintani | Nucleon studies for the standard model and beyond â€“ Measurement of neutron EDM and proton decay, and Strangeness in the Nucleon with Domain Wall fermions |
| Robert Sugar | QCD with Four Flavors of Highly Improved Staggered Quarks fermions |
| Oliver Witzel | B-meson physics with domain-wall light quarks and relativistic heavy quarks fermions/ |

**2012 Type B Proposals (less than 2,500,000 J/Psi core-hours)**

| PI | Project |
|---|---|
| Andrei Alexandru | Sea quark effects in hadron electric polarizability |
| Simon Catterall | Simulating N = 4 Super Yang-Mills using GPUs |
| Michael Engelhardt | Electric spin polarizability of the neutron |
| Michael Engelhardt | Nucleon transverse momentum dependent parton distribution functions with domain wall fermions on fine lattices |
| Joel Giedt | Studies of the infrared fixed point in SU(2) gauge theory with two flavors of adjoint fermions |
| Tomomi Ishikawa | Application of low-mode averaging to B0 â˘ B0bar mixing with static heavy quark and domain-wall light quarks |
| Yu Maezawa | Meson screening masses at finite temperature with Highly Improved Staggered Quarks |
| Dhagash Mehta | Transport Coeffcients from Lattice QCD with 2 Dynamical Fermions |
| James OsbornÂ Â Â | Disconnected contributions to nucleon form factors with chiral fermions |
| Dru Renner | Step-scaling methods for operator mixing |
| Stephen Sharpe | Non-perturbative renormalization with improved staggered fermions |